

EXHIBIT G

From: Paul King
To: steve@stelevachani.com
Subject: Re: more on orkut...
Date: Monday, September 12, 2005 12:41:51 AM

I don't know how easy rotating IP addresses is.

It is certainly possible. I know it is fairly easy to make one machine accept many incoming IPs. What I don't know is how you rotate the outgoing IPs (although I'm sure it is possible). Someone who knows linux network driver configuration would know the answer.

If you had to reboot the server in between IP changes, it could be a slow process. And if you can't use multiple IP addresses in parallel, the could be much slower.

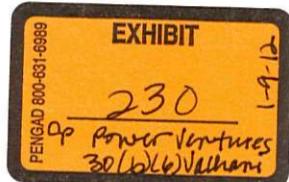
On this last point, the reason is this: It takes Orkut about 1 sec to generate a web page. That would be 60 pages per second if you did them in sequence. So you have to parallelize it, and have 100+ page requests in process simultaneously. If you can only use one IP at a time, then you have to switch IPs once-per-minute. If you have to reboot to change an IP address (hopefully not!), then the requesting server might need to reboot once per minute or once per 10 minutes.

Paul

----- Original Message -----

From: "Steven Vachani" <steve@stelevachani.com>
To: "Paul King" <paul@pkng.org>; <steve@stelevachani.com>
Sent: Monday, September 12, 2005 12:31 AM
Subject: Re: more on orkut...

> One of our plans it to have rotating IP addresses. So
> we could have 10,000 or even more IP addresses
> throughout this process.
>
> I assume the rotating IP addresses is something easy.
> We could obtain 25,000 IP addresses for example.
>
> What do you think?
>
> --- Paul King <paul@pkng.org> wrote:
>
> > Steve,
> >
> > An orkut profile page, not counting the photos, is
> > about 20k of HTML. Of
> > this, there appears to be about 1K of actual data.
> >
> > For the friends data, each "Friends" page displays
> > 18 friends using about
> > 20k of HTML which contians 80 bytes of data, which
> > expands to 400 bytes of
> > data in a database. If the "average" person has 180
> > friends, this is 10
> > friends pages, that's another 200k transferred for



> > 4k of friend data in the
> > database.
> >
> > Communities display at 10 per page at 20k each
> > (because of all the
> > description info). So 100 communities is another 10
> > pages 40 bytes of data
> > per page or 200 bytes in a database. If the average
> > person has 100
> > communities, that is 10 community pages for a total
> > of 200k transferred, and
> > 2k stored in the database.
> >
> > So the net of it for a single profile, friends, and
> > communities is
> > approximately:
> > 20 transferred pages (1 profile, 10 friends, 10
> > communities)
> > 400k transferred HTML
> > 7k in the database
> >
> > Downloading 10M users would be:
> > 200M transferred pages
> > 4000 GB transferred HTML
> > 70 GB stored in the database
> >
> > You would need the bandwidth for the transferred
> > HTML. And you would not
> > want to overwhelm the orkut servers, which are
> > pretty slow anyway.
> >
> > One server running continuously should be able to
> > download 1000 pages per
> > minute, if Orkut can keep up. 10 servers would be
> > 10k pages per minute.
> > With 10 servers running in parallel, you would need
> > 300 hours (14 days) of
> > continuous non-stop operation to get all the text
> > data.
> >
> > The good news is that 1M users should only take 1 -
> > 2 days, or only need one
> > server for 2 weeks.
> >
> > There may be ways of making the querying servers go
> > faster, but Orkut may
> > impose limits on the number of page requests from a
> > single IP address.
> >
> > Orkut's license agreement explicitly forbids robots,
> > so they are onto it and
> > may be doing things to prevent automated retrieval.
> >
> > Paul
> >
> >
>